Hiding-Based Compression for Improved Color Image Coding

Patrizio Campisi[a], Deepa Kundur[b], Dimitrios Hatzinakos[b], and Alessandro Neri[a]

[a] Dip. Ingegneria Elettronica, Università degli Studi di Roma "Roma Tre",
via Della Vasca Navale 84, I-00146 Roma, Italy
[b]Edward S. Rogers Sr. Dept. of Electrical and Computer Engineering, University of Toronto,
Toronto, Ontario, Canada M5S 3G4

ABSTRACT

This paper considers the use of data hiding strategies for improved color image compression. Specifically, color information is "piggybacked" on the luminance component of the image in order to reduce the overall signal storage requirements. A practical wavelet-based data hiding scheme is proposed in which selected perceptually irrelevant luminance bands are replaced with perceptually salient chrominance components. Simulation results demonstrate the improvement in compression quality of the proposed scheme to SPIHT and JPEG at low bit rates. The novel technique also has the advantage that it can be used to further reduce the storage requirements of algorithms such as SPIHT which is optimized for grayscale image compression.

Keywords: multiresolution analysis, compression, perceptual coding, color image processing, data hiding, wavelet-based watermarking.

1. INTRODUCTION

Data hiding within multimedia has received growing interest in recent years due to its potential for signal captioning, maintaining audit trails in media commerce, and copy protection through the development of digital watermarking technology. By embedding key information within the media itself, it is safe from content separation. Data hiding is the general process by which a discrete information stream is merged within media content by imposing imperceptible changes on the original *host* signal.

One of the main obstacles within the data hiding community has been developing a scheme which is robust to *perceptual coding*. *Perceptual coding* refers to the *lossy* compression of multimedia signal data using human perceptual models; the compression mechanism is based on the premise that minor modifications of the signal representation will not be noticeable in the displayed signal content. These modifications are imposed on the signal in such a way as to reduce the number of information bits required for storage of the content. Human perceptual models are often theoretically and experimentally derived to determine the changes on a signal which remain imperceptible. A duality exists between the problems of perceptual coding and data hiding; the former problem attempts to remove irrelevant and redundant information from a signal, while the latter uses the irrelevant information to mask the presence of the hidden data. Thus, the goals of data hiding and perceptual coding can be viewed as being somewhat contradictory. As a result, several papers have dealt with integrating perceptual coding with data hiding[1–7] and others have investigated the theoretical relationship between both processes.[8]

The central theme of all the works cited above is that there must be an appropriate *compromise* between data hiding and compression to develop a method which performs both reasonably. It is assumed that each process hinders, not helps, the objective of the other. Specifically, data hiding decreases the overall possible compression ratio and perceptual coding tampers with the hidden information, so that extraction is difficult.

In this paper, we take a different perspective. We demonstrate how judicious data hiding can be used for improved practical compression. In instances where a specific web browser or program for dynamic information

---

Further author information: (Send correspondence to P.C.)
P.C.: E-mail: campisi@uniroma3.it, Telephone: +39 06 5517 7064
D.K.: E-mail: deepa@comm.toronto.edu, D.H.: E-mail: dimitris@comm.toronto.edu, A.N.: neri@uniroma3.it

Security and Watermarking of Multimedia Contents IV, Edward J. Delp III, Ping Wah Wong,
Editors, Proceedings of SPIE Vol. 4675 (2002) © 2002 SPIE · 0277-786X/02/$15.00

swapping is used, the application may support only a limited number of image compression algorithms. It is valuable to be able to improve the quality of the compressed information transparently. One approach would be to preprocess the signal in order to improve upon potential compression artifacts. In this paper, we employ robust data hiding at this preprocessing stage in order to pass or tag signal information reliably through the compression stage. The information then can be used, after compression, in order to improve signal fidelity or provide some other form of added-value.

As a start, we address the case of color image compression; we "piggy-back" chrominance information on the luminance component of the image to improve the signal fidelity upon post-compression reconstruction for a fixed compression ratio. We make use of the popular wavelet-based compression algorithm SPIHT,[9] so that data hiding can also be effectively accomplished in this multiresolution-like domain.

Our novel technique essentially transforms a given color image into the YIQ color space where the chrominance information is sub-sampled and embedded in the selected components of the wavelet domain of the luminance component. The embedding strategy as well as the subbands employed for data hiding provide an effective trade-off between information capacity and robustness. The resulting signal is then passed through the SPIHT compression algorithm.

We demonstrate how our technique can be used as preprocessing to improve the performance of popular image compression schemes such as SPIHT that are optimized for grayscale image compression. Simulation results demonstrate the superior performance of the proposed technique in comparison to JPEG and straightforward SPIHT. For low rates such as 0.15-0.45 bpp our proposed hiding-based preprocess provides significantly improved perceptual results.

1.1. Objectives of this Paper

In this work we present an approach to improve the efficiency of compression by incorporating data hiding principles. On a larger scale, the presented work aims to, in part, investigate the contradictory processes of data hiding and compression in order to derive insights into effective means to merge them.

Specifically, we wish

- to develop a simple approach to enhance existing compression algorithms; instead of using a coarser quantization table, we propose to incorporate data hiding to achieve this goal.
- to design a compression scheme in which color information is "piggybacked" on the grayscale component to provide the option of viewing the information in color or as a monochrome signal.
- to compare our proposed data hiding-based compression approach practically with SPIHT, an effective wavelet-based compression algorithm.

2. HYPOTHESIS

Consider the signal $f_0(x)$ which represents an audio signal, image or video sequence. There exists a family of functions in the set $\mathbf{P}(f_0)$ which are perceptually identical to $f_0(x)$. Thus, if $f_k(x) \in \mathbf{P}(f_0)$ then we know that $f_k(x)$ is perceptually identical to $f_0(x)$.

In ideal compression every possible perceptually identical signal is mapped to the same representation. Thus, all signals in the set $\mathbf{P}(f_0)$ will be collapsed into one compressed signal. In data hiding, information is embedded into a host signal $f_0(x)$ by modifying it so that the resulting signal is perceptually identical; therefore this new signal is also in the set $\mathbf{P}(f_0)$. It then follows that ideal compression applied to a signal containing hidden information has the effect of annihilating the discreet data.

However, in practice, compression is not completely efficient; that is, there exists some irrelevant information which has not been removed. The non-ideality comes from the constraint that the coder have structure, and the use of inadequate perceptual models to account for masking characteristics. In terms of our description above, for practical compression not all signals in $\mathbf{P}(f_0)$ are mapped to the same representation. Thus, there is a small bandwidth available for data hiding. If this bandwidth could be used to transmit information about the signal such as chrominance components then even greater practical compression may be achieved. This approach cannot help improve the compression ratio in the case of ideal perceptual coding, but can improve the situation in the case of practical compression.

The approach can potentially provide improved performance when we are restricted to use a coding scheme which is not very efficient. For example, one of the most popular wavelet image compression schemes SPIHT is optimized for grayscale compression, and may not perform acceptably for color images. If the technique outlined above could be applied to “piggy-back” the chrominance components in a compressed version of the luminance image, then some performance advantage may be established. Thus, we propose our scheme as a possible improvement to existing approaches which may not performance optimally. The also has the advantage in that the color chrominance is embedded in the grayscale, so that it may be viewed later even if color is not initially important.

3. DATA HIDING-BASED COMPRESSION

The general principle behind the proposed color image hiding-based compression scheme is discussed in this section. With reference to Fig. 1, let us consider a color image $X[n_1, n_2]$.

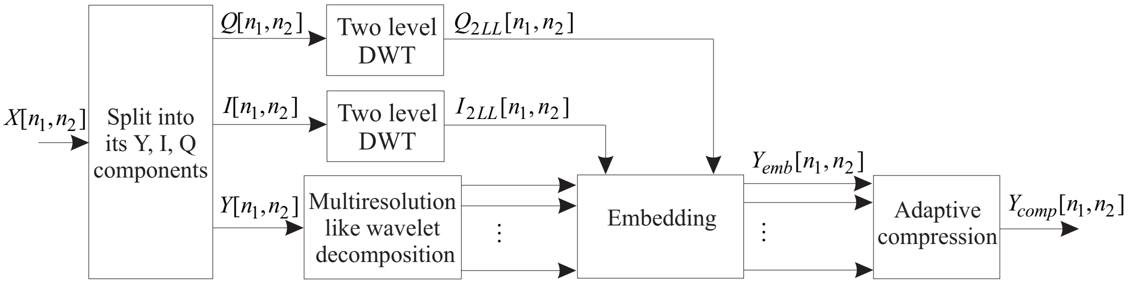

Figure 1: *Compressive Data Hiding Scheme for Color Images.*

First, the given image is split into its three components in the YIQ color space, where $Y[n_1, n_2]$ is the luminance and $I[n_1, n_2]$ and $Q[n_1, n_2]$ jointly represent hue and saturation of the color.

It is worth noting that the choice of the color space where to perform the decomposition of the given color image is relevant to the problem we have addressed. In fact, since our goal is compression, color spaces, such as the RGB space, where there is a significant correlation between the three color components should be avoided. On the contrary, the color spaces YIQ, YUV and YC_bC_r nearly provide as much energy compaction as a theoretically optimal decomposition performed using the Karhunen-Loeve transform. Since they are equivalent for our application, we have resorted to split the given color image into its three components in theYIQ color space, where the Y coordinate represents the luminance $Y[n_1, n_2]$ and the I and Q coordinates represent the chrominance components $I[n_1, n_2]$ and $Q[n_1, n_2]$ respectively. The I and Q components are used to jointly represent saturation and hue.

The bandwidths of the last two components are much smaller than that of the luminance one. This implies that the chrominance components can be represented by their subsampled versions without any perceptual loss of quality for the reconstructed image.

Second, we hide or “piggyback” the color data in the luminance component such that the perceptual appearance of the luminance component of the host image is not impaired. The embedding is performed in the wavelet domain since the wavelet transform can provide both spatial and frequency localization which suits the behavior of the human visual system (HVS).[10, 11] By taking into account the HVS, we can judiciously select the regions of the luminance subbands that can transparently hide the chrominance components effectively.

More specifically, an unconventional two-level multiresolution-like wavelet decomposition is performed on the luminance component. With reference to Fig. 2, the first level of the multiresolution decomposition is obtained by performing a discrete wavelet transform (DWT) onto $Y[n_1, n_2]$. In particular an 8 tap Daubachies filter is

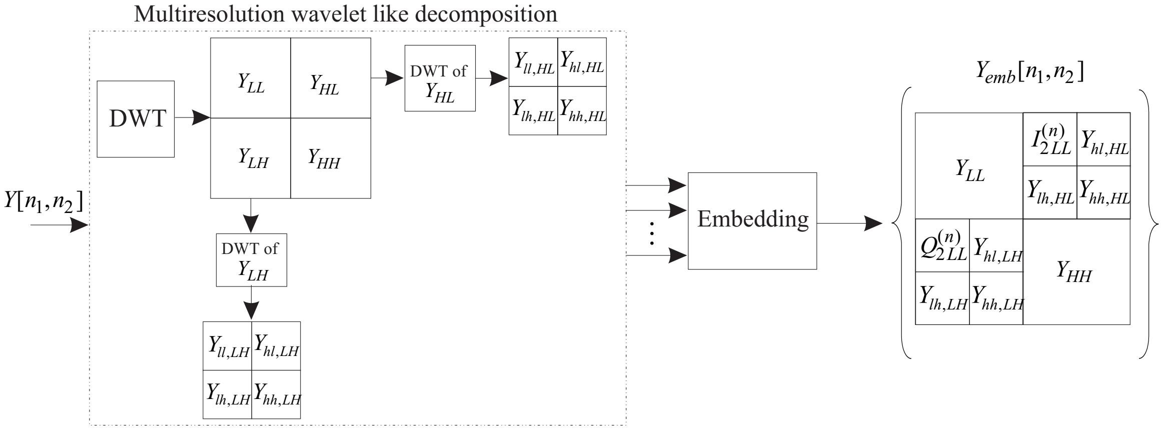

Figure 2: *Multiresolution-like wavelet decomposition and subbands structure resulting from the embedding process.*

used. This leads to the subbands $Y_{LL}[n_1, n_2]$ and $Y_{LH}[n_1, n_2]$, $Y_{HL}[n_1, n_2]$, and $Y_{HH}[n_1, n_2]$, which take into account the image at coarser resolution plus the “horizontal”, “vertical”, and “diagonal” details of the image also at coarser resolution, respectively.

The subbands $Y_{LH}[n_1, n_2]$ and $Y_{HL}[n_1, n_2]$ are chosen to host the chrominance information. The rational behind this choice relies on the observation that, in order to obtain a good trade-off between robustness and transparency, many watermarking techniques (e.g. see Refs. 12, 13 and references therein) use “middle frequency” coefficients. This makes the subbands $Y_{LH}[n_1, n_2]$ and $Y_{HL}[n_1, n_2]$ suitable to host the data, whereas the subband $Y_{HH}[n_1, n_2]$ is not.

Then the subbands $Y_{LH}[n_1, n_2]$, $Y_{HL}[n_1, n_2]$ are further wavelet decomposed thus leading to the subbands $Y_{\alpha,\beta}[n_1, n_2]$ with $\alpha \in \{ll, hl, lh, hh\}$ and $\beta \in \{HL, LH\}$ [*]. In particular, $Y_{ll,HL}[n_1, n_2]$ and $Y_{ll,LH}[n_1, n_2]$ represent the low-pass subbands, at coarser resolution, obtained from the high frequency subbands $Y_{HL}[n_1, n_2]$ and $Y_{LH}[n_1, n_2]$ respectively. It is expected that their energy contribution is relatively small compared to the energy of the remaining subbands of the set $Y_{\alpha,\beta}[n_1, n_2]$. Thus, they can be neglected without impairing the quality of the reconstructed image.

This conjecture has been experimentally verified upon a wide range of images’ typology using both subjective evaluation criteria such as visual perceptibility measures and objective ones like the Peak Signal to Noise Ratio (PSNR), given by

$$PSNR = 10 \log_{10}\left(\frac{255^2}{MSE}\right) \quad [dB]$$

where

$$MSE = \frac{1}{N^2} \sum_{n_1=0}^{N-1} \sum_{n_2=0}^{N-1} \left(Y[n_1, n_2] - Y_{mod}[n_1, n_2]\right)^2$$

is the mean square error between the original image $Y[n_1, n_2]$, of NxN pixels, and a modified replica $Y_{mod}[n_1, n_2]$. To assess the general perceptual irrelevance of the $Y_{ll,HL}[n_1, n_2]$ and $Y_{ll,LH}[n_1, n_2]$ bands, we consider a series

---

[*] This is the *unconventional* part of the scheme because usually the $Y_{LL}[n_1, n_2]$ band is further decomposed instead of its details. However, our proposed decomposition is necessary for an appropriate compromise between imperceptibility and robustness of our embedding stage as we will discuss later in the section.

	$PSNR_{zeroed}$	$PSNR_{embed}$
Baboon	33.4	32.1
Lena	36.2	35.3

Table 1: Peak SNR.

of test images with widely varying characteristics. For each image, a modified replica is produced by zeroing only the subbands $Y_{ll,HL}[n_1, n_2]$ and $Y_{ll,LH}[n_1, n_2]$ and keeping the remaining perfectly intact. The PSNRs of the resulting images (denoted by $PSNR_{zeroed}$) are presented in the first column of Table 1. The values are reasonably high; in addition, there is no perceptual change in the quality of the modified signal $Y_{mod}[n_1, n_2]$ in each test case. Therefore, these perceptually negligible bands can be zeroed and replaced with the chrominance information.

However, the chrominance bands need to be preprocessed before the embedding in order to obtain a parsimonious representation. Furthermore, they undergo a two level wavelet decomposition. Since, as already outlined, their bandwidth is much smaller than the luminance's one, the subbands $I_{2LL}[n_1, n_2]$ and $Q_{2LL}[n_1, n_2]$, that is the lowpass chrominance replicas at the coarsest resolution of the performed pyramidal decomposition, can be used to reconstruct the color image without any perceptual quality loss.

Therefore we choose to embed $I_{2LL}[n_1, n_2]$ and $Q_{2LL}[n_1, n_2]$ into the subbands $Y_{ll,HL}[n_1, n_2]$ and $Y_{ll,LH}[n_1, n_2]$ which have been previously zeroed, thus obtaining $Y_{emb}[n_1, n_2]$ (see Fig. 2). However, before the embedding, the energies of $I_{2LL}[n_1, n_2]$ and $Q_{2LL}[n_1, n_2]$ have to be normalized to the values of the corresponding host subbands as not to impair the perceptual appearance of the reconstructed image. It should be noted that the normalization values have to be transmitted to the decoder since they are necessary to reconstruct properly the color information. To this end they can be embedded in the header of the image.

We have experimentally verified that the aforementioned embedding procedure causes no perceptual degradation to the luminance component of the signal. Some of these results are reported in Fig. 3 where $Y_{emb}[n_1, n_2]$, the luminance image component with the embedded chrominance information, is shown for different test cases.

It can be observed that the luminance after the embedding appears perceptually indistinguishable from the host image. Also a quantitative evaluation is performed by calculating the PSNR, denoted $PSNR_{embed}$, whose values for different test images, are presented in the second column of Table 1.

The final step of the algorithm consists of the *adaptive compression stage* (see Fig. 1) that codes each of the $Y_{emb}[n_1, n_2]$ "subbands" separately in order to guarantee the desired bit rate, related to the maximum distortion allowed, for each subband. Specifically, a wavelet-based coder is used instead of a DCT-based coder; it has been proven that wavelet based coders provide better rate-distortion performance than the DCT-based JPEG and also allow a progressive coding approach.[14] The method of set partitioning in hierarchical tree (SPIHT) has been employed in the proposed scheme.

At the receiving side, the color image is reconstructed, from the compressed bit stream, performing dual operations of the ones done during the coding stage. First the single subbands are decoded, thus obtaining an estimation (denoted by $\widehat{\ }$) $\widehat{Y}_{emb}[n_1, n_2]$ of $Y_{emb}[n_1, n_2]$. Then $\widehat{I}_{2LL}[n_1, n_2]$ and $\widehat{Q}_{2LL}[n_1, n_2]$ are extracted and the components $\widehat{I}[n_1, n_2]$ and $\widehat{Q}[n_1, n_2]$ are obtained by applying a two level inverse DWT. After having zeroed $\widehat{Y}_{ll,HL}[n_1, n_2]$ and $\widehat{Y}_{ll,LH}[n_1, n_2]$, the subbands $\widehat{Y}_{HL}[n_1, n_2]$ and $\widehat{Y}_{LH}[n_1, n_2]$ are reconstructed by performing a one level inverse DWT. Finally an estimation of the luminance $\widehat{Y}[n_1, n_2]$ is achieved.

4. SIMULATION RESULTS

In this section the effectiveness of the proposed method is discussed. Experimentations have been performed on color images represented by 24 bpp, whose luminance components (Y) and whose chrominance components (I and Q) are displayed in Figs. 4, 5 (luminance), in Figs. 6, 7 (chrominance I) and in Figs. 8, 9 (chrominance Q) respectively.

Figure 3. *First row: left: original grayscale Baboon, right: Baboon with the chrominance components I and Q embedded. Second row: left: original grayscale Lena, right: Lena with the chrominance components I and Q embedded.*

In the same figures the corresponding compressed replicas obtained using the proposed approach are shown. Moreover for the sake of comparison, are also shown the compressed version of the original images obtained using JPEG and SPIHT. The color images have been compressed at (0.15 bpp, and 0.45 bpp). It is worth pointing out that at the bit rate 0.15 bpp the JPEG compressed images have been compressed at 0.25 bpp, instead of 0.15 bpp, since the JPEG coder does not allow further compression. It is evident that the proposed method performs better than JPEG and outperforms SPIHT that is not optimized for color images.

The assessment of the performance of the proposed method requires the quantification of the perceptual error between two color images. To this end we resort to use the perceptive uniform color space L^* a^* b^* and to employ as metric the normalized color distance (NCD)[15] defined as follows

$$\mathrm{NCD} = \frac{\sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} \Delta E[n_1, n_2]}{\sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} E[n_1, n_2]} \tag{1}$$

where $\Delta E[n_1, n_2]$ is given by

$$\Delta E[n_1, n_2] = \sqrt{(L_o^*[n_1, n_2] - L_c^*[n_1, n_2])^2 + (a_o^*[n_1, n_2] - a_c^*[n_1, n_2])^2 + (b_o^*[n_1, n_2] - b_c^*[n_1, n_2])^2} \tag{2}$$

having indicated with the sub indices "o" and "c" the color components of the original and compressed image respectively and with $E[n_1, n_2] = \sqrt{(L_o^*[n_1, n_2])^2 + (a_o^*[n_1, n_2])^2 + (b_o^*[n_1, n_2])^2}$ the Euclidean norm of the pixel in position $[n_1, n_2]$ belonging to the uncompressed image.

This quantitative performance evaluation, at different bit rates (0.15 bpp and 0.45 bpp), is performed on the images obtained by applying our compressive data hiding approach, JPEG and SPIHT with respect to the original image and the results are shown in Table 2.

In this way, our scheme, as outlined in Section 2, can be seen as possible improvement for methods that do not perform optimally for color images.

Figure 4. *First row: Baboon, second row: Lena. From left to right: luminance component of the original image (original color image 24 bpp), compressed luminance component using the proposed approach (resulting color image 0.15 bpp), compressed luminance component using the SPIHT method (resulting color image 0.15 bpp), compressed luminance component image using JPEG (resulting color image 0.25 bpp).*

Figure 5. *First row: Baboon, second row: Lena. From left to right: luminance component of the original image (original color image 24 bpp), compressed luminance component using the proposed approach (resulting color image 0.45 bpp), compressed luminance component using the SPIHT method (resulting color image 0.45 bpp), compressed luminance component image using JPEG (resulting color image 0.45 bpp).*

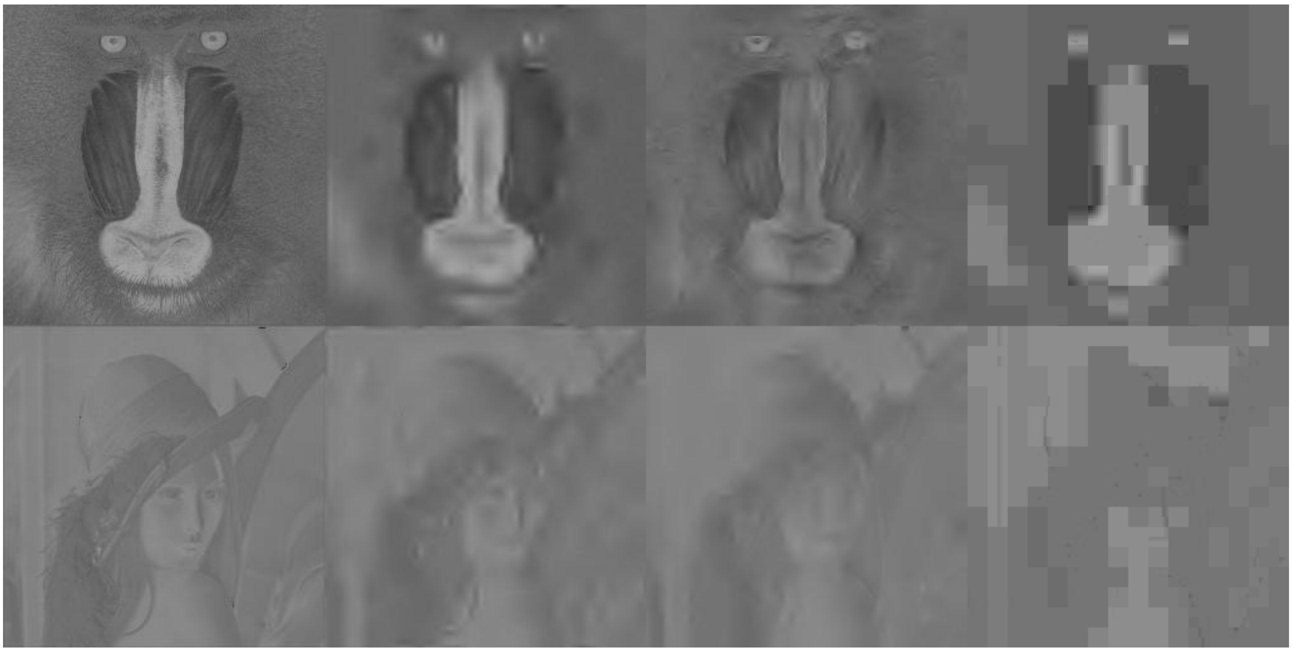

Figure 6. *First row: Baboon, second row: Lena. From left to right: chrominance I of the original image (color image 24bpp), compressed chrominance I using the proposed approach (resulting color image 0.15 bpp), compressed chrominance I using the SPIHT method (resulting color image 0.15 bpp), compressed chrominance I using JPEG (resulting color image 0.25 bpp).*

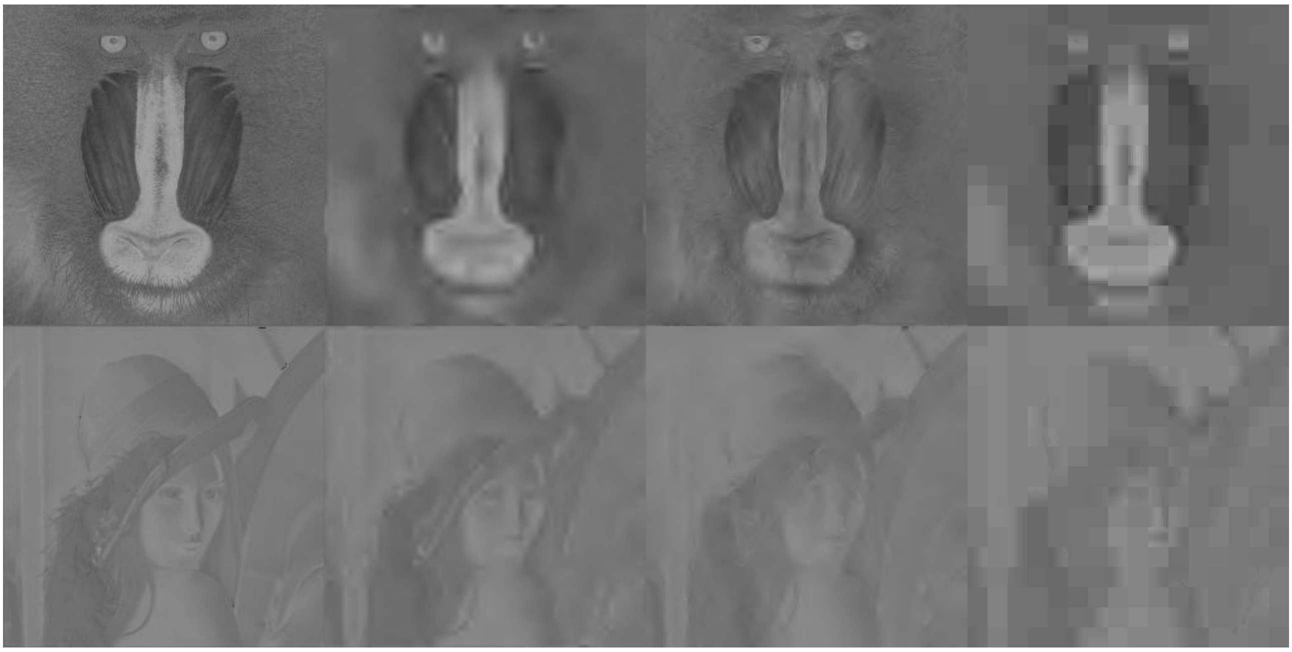

Figure 7. *First row: Baboon, second row: Lena. From left to right: chrominance I of the original image (color image 24bpp), compressed chrominance I using the proposed approach (resulting color image 0.45 bpp), compressed chrominance I using the SPIHT method (resulting color image 0.45 bpp), compressed chrominance I using JPEG (resulting color image 0.45 bpp).*

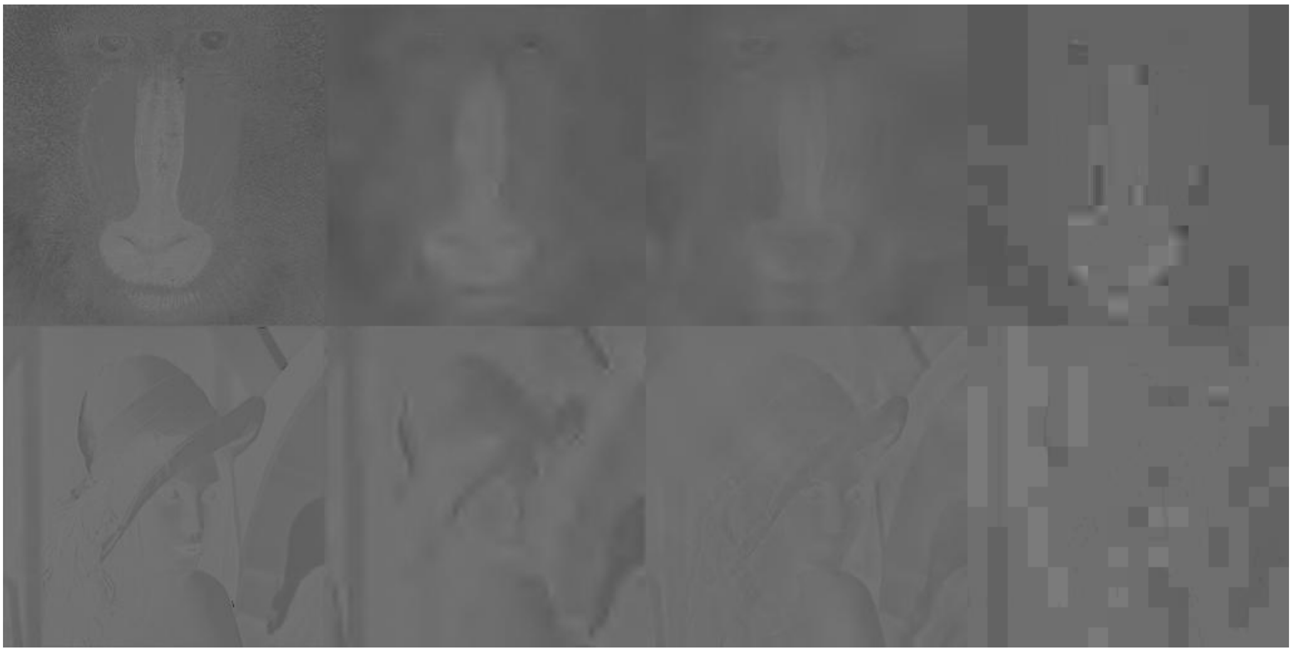

Figure 8. *First row: Baboon, second row: Lena. From left to right: chrominance Q of the original image (color image 24bpp), compressed chrominance Q using the proposed approach (resulting color image 0.15 bpp), compressed chrominance Q using the SPIHT method (resulting color image 0.15 bpp), compressed chrominance Q using JPEG (resulting color image 0.25 bpp).*

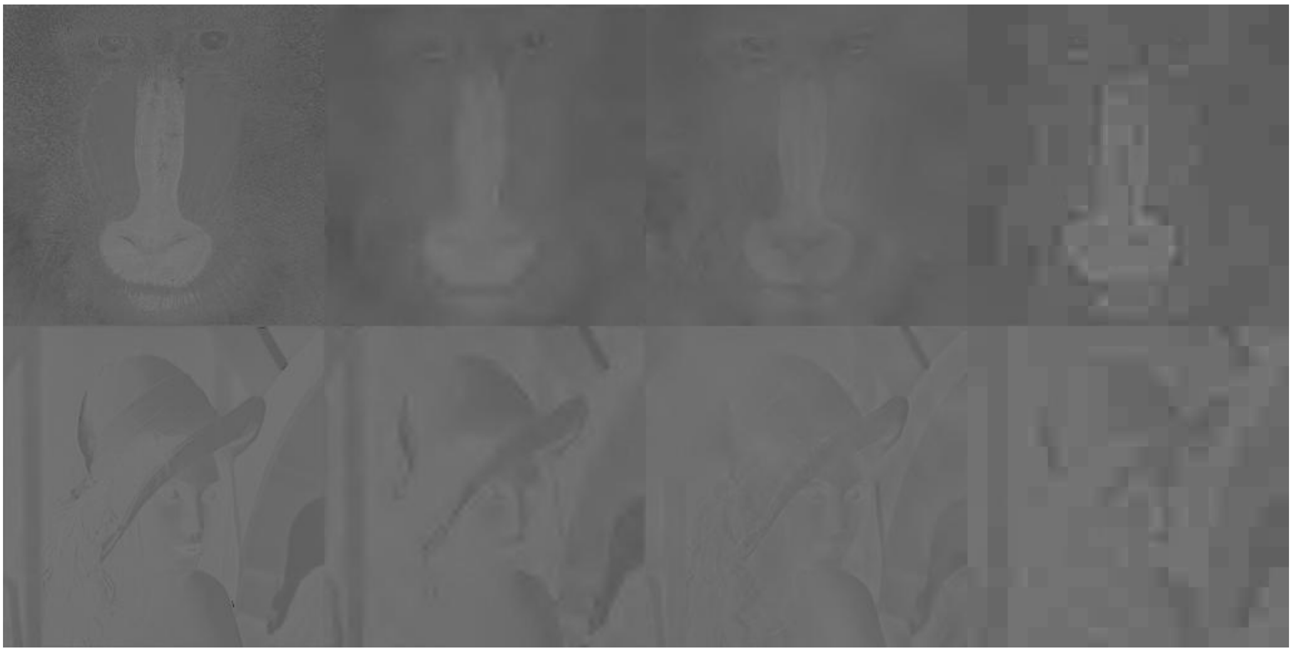

Figure 9. *First row: Baboon, second row: Lena. From left to right: chrominance Q of the original image (color image 24bpp), compressed chrominance Q using the proposed approach (resulting color image 0.45 bpp), compressed chrominance Q using the SPIHT method (resulting color image 0.45 bpp), compressed chrominance Q using JPEG (resulting color image 0.45 bpp).*

	bit rate	Data hiding	SPIHT	JPEG
Baboon	0.15	0.1492	0.1985	0.19
	0.45	0.1225	0.1896	0.1545
Lena	0.15	0.0932	0.1650	0.1753
	0.45	0.0695	0.1563	0.0881

Table 2: NCD evaluation for the compression rates 0.15 bpps, 0.45 bpps.

5. FINAL REMARKS

In this paper, a data hiding-based color image coding scheme is proposed in which the process of data embedding is used to enhance instead of hinder the compression process. The method involves three basic stages. First, an unconventional two level wavelet decomposition is performed on the luminance component of the color image; the perceptually irrelevant subbands are selected. Second, these subbands are replaced with the perceptually salient components of the chrominance information. Last, the resulting signal is then compressed using an algorithm such as SPIHT. The result is a grayscale coded image in which the color information is "piggybacked" without impairing the overall perceptual quality of the composite image. This representation provides versatility in viewing the information in either grayscale or color. Furthermore, it may be considered a preprocessing stage for color images in order to make them more suitable to grayscale optimized image coding algorithms. Through simulations we show that, for a given bit rate, data hiding-based compression achieves better perceptual quality with respect to well consolidated coding schemes such as JPEG and SPIHT at low bit rates.

REFERENCES

1. J. Lacy, S. R. Quackenbush, A. R. Reibman, D. Shur, and J. H. Snyder, "On Combining Watermarking with Perceptual Coding", *Proc. Int. Conf. on Acoustics, Speech and Signal Processing*, Vol. 6, pp.3725-3728, May 1998.
2. H.-J. Wang and C.-C. J. Kuo, "An Integrated Progressive Image Coding and Watermark System", *Proc. Int. Conf. on Acoustics, Speech and Signal Processing*, Vol. 6, pp.3721-3724, May 1998.
3. J. Meng and S.-F. Chang, "Embedding Visible Video Watermarks in the Compressed Domain", *Proc. IEEE Int. Conference in Image Processing*, Vol. 1, October 1998.
4. S. Bhattacharjee and M. Kutter, "Compression Tolerant Image Authentication", *Proc. IEEE Int. Conference in Image Processing*, Vol. 1, October 1998.
5. T.-Y. Chung, M.-S. Hong, Y.-N. Oh, D.-H. Shin and S.-H. Park, "Digital Watermarking for Copyright Protection of MPEG2 Compressed Video", *IEEE Transactions on Consumer Electronics*, Vol. 44, No. 3, pp.895-901, August 1998.
6. J. Lacy, S. R. Quackenbush, A. R. Reibman, and J. H. Snyder "Intellectual Property Protection Systems and Digital Watermarking", *Optics Express*, Vol. 3, No. 12, pp.478-484, December 1998.
7. D. Kundur and D. Hatzinakos "Mismatching Perceptual Models for Effective Watermarking in the Presence of Compression", *Proc. SPIE, Multimedia Systems and Applications II*, September 1998.
8. D. Kundur, "Energy Allocation for High Capacity Watermarking in the Presence of Compression", *Proc. IEEE Int. Conf. on Image Processing*, September 2000.
9. A. Said, W.A. Pearlman, "A New Fast and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees", *IEEE Trans. on Circuits and System for Video Technology*, Vol. 6, No. 3, pp.243-250, June 1996.
10. S. Mallat, "Multifrequency Channel Decompositions of Images and Wavelet Models," *IEEE Transactions on Acoustic, Speech and Signal Processing*, vol. 37, No. 12., December 1989.
11. S. Mallat, "Wavelet for a vision", *Proceedings of the IEEE*, Vol. 84, No.4, pp. 604-614, April 1996.
12. I.J. Cox, J. Kilian, T. Leighton and T. Shamoon, "Secure Spread Spectrum Watermarking for Multimedia," *IEEE Trans. on Image Processing*, vol. 6, no. 12, pp. 1673-1687, December 1997.
13. F. Hartung and M. Kutter "Multimedia Watermarking Techniques," *Proceedings of the IEEE*, vol. 87, no. 7, pp. 1079-1107, July 1999.
14. Jin Li, Po-Yuen Cheng, and C.-C Jay Kuo, "A Wavelet Transform Approach to Video Compression", *Proc. SPIE Wavelet Applications II*, Orlando, Fl, April 1995.
15. K. N. Plataniotis, A. N. Venetsanopoulos, **Color Image Processing and Applications**, Springer 2000.